Sharing Information Without Regret in Managed Stochastic Games

Michael J. Neely
University of Southern California
http://www-bcf.usc.edu/~mjneely

arXiv:1412.8736v1 [math.OC] 30 Dec 2014

***Abstract*— This paper considers information sharing in a multi-player repeated game. Every round, each player observes a subset of components of a random vector and then takes a control action. The utility earned by each player depends on the full random vector and on the actions of others. An example is a game where different rewards are placed over multiple locations, each player only knows the rewards in a subset of the locations, and players compete to collect the rewards. Sharing information can help others, but can also increase competition for desirable locations. Standard Nash equilibrium and correlated equilibrium concepts are inadequate in this scenario. Instead, this paper develops an algorithm where, every round, all players pass their information and intended actions to a game manager. The manager provides suggested actions for each player that, if taken, maximize a concave function of average utilities subject to the constraint that each player gets an average utility no worse than it would get without sharing. The algorithm acts online using information given at each round and does not require a specific model of random events or player actions. Thus, the analytical results of this paper apply in non-ergodic situations with any sequence of actions taken by human players.**

I. Introduction

This paper considers a stochastic game where each player has incomplete information. A central issue is whether or not this information should be shared. Indeed, players with special access to desirable information may prefer to keep this information private. The goal of this paper is to design an efficient collaborative strategy that allows players to share information without sacrificing their own interests.

The general game structure is as follows: There are N players that repeatedly play a game over a sequence of rounds $t \in \{0, 1, 2, \ldots\}$. On each round t, there is a random event vector $\omega(t) = (\omega_1(t), \ldots, \omega_M(t))$ that describes characteristics of the game for that round. The value M is a positive integer that can be different from N. Each player can observe a *portion* of the components of the $\omega(t)$ vector. Specifically, for each $i \in \{1, \ldots, N\}$, define $\mathcal{S}_i$ as the subset of indices in $\{1, \ldots, M\}$ that are observable by player i. Every round t, each player i observes its components of $\omega(t)$ and then chooses an action $\alpha_i(t)$ based on this (incomplete) information. Define $\alpha(t) = (\alpha_1(t), \ldots, \alpha_N(t))$ as the joint action vector. The resulting round-t payoff for player i is $u_i(t)$, also called the *utility*. The utility $u_i(t)$ is a general function of $\alpha(t)$ and $\omega(t)$:

$$u_i(t) = \hat{u}_i(\alpha(t), \omega(t)) \;\; \forall i \in \{1, \ldots, N\}$$

The author is with the Electrical Engineering department at the University of Southern California, Los Angeles, CA.

This work is supported in part by the NSF Career grant CCF-0747525.

Each player wants to earn a large time average utility $\overline{u}_i$:

$$\overline{u}_i = \lim_{t\to\infty} \frac{1}{t} \sum_{\tau=0}^{t-1} u_i(\tau)$$

Pooling information about the $\omega(t)$ vector and making a *team decision* can improve the sum utility. However, individual players may want to keep their information private to increase their own utility.

A. Example game structure

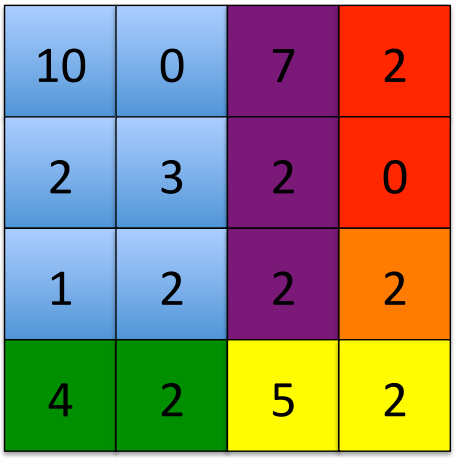

Fig. 1. An illustration of a 3-player location-reward game. Locations known only to player i are blue, yellow, red for $i \in \{1, 2, 3\}$, respectively. Locations known to both players 1 and 2 are green, players 2 and 3 are orange, and players 3 and 1 are purple.

Consider a square region that is partitioned into M disjoint sub-regions, called *locations* (see Fig. 1 with $M = 16$ locations). Every round, a random reward $\omega_m(t)$ appears in each location $m \in \{1, \ldots, M\}$. Let $\omega(t) = (\omega_1(t), \ldots, \omega_M(t))$ be the vector of current rewards. For example, $\omega(t)$ might be a random vector that is independent and identically distributed (i.i.d.) over rounds $t \in \{0, 1, 2, \ldots\}$ with some arbitrary joint probability distribution. Suppose there are 3 players:

- Player 1 knows rewards for the blue, purple, and green squares.
- Player 2 knows rewards for the yellow, green, and orange squares.
- Player 3 knows rewards for the red, orange, and purple squares.

Every round, each player chooses a single location where it competes for the current reward. Specifically, for each $i \in \{1,2,3\}$, let $\mathcal{A}_i$ be a subset of $\{1,\ldots,M\}$ that represents the set of locations player i is allowed to choose from, called the *action set* for player i. The sets $\mathcal{A}_i$ and $\mathcal{S}_i$ can be different, so that a player might choose a location in which she does not know the reward. Let $\alpha_i(t)$ be the location in $\mathcal{A}_i$ chosen by player i on round t. If a player is the only one to choose a certain location m, she earns the full reward $\omega_m(t)$. Else, the reward is split evenly amongst all players who choose that location. Specifically, for each $m \in \{1,\ldots,M\}$ define $K_m(t)$ as the number of players who choose location m on round t. The resulting utility for player $i \in \{1,2,3\}$ is:

$$u_i(t) = \frac{\omega_{\alpha_i(t)}(t)}{K_{\alpha_i(t)}(t)}$$

This utility is indeed a function of the vectors $\alpha(t)$ and $\omega(t)$:

$$\hat{u}_i(\alpha,\omega) = \frac{\omega_{\alpha_i}}{\sum_{n=1}^{3} 1\{\alpha_n = \alpha_i\}} \tag{1}$$

where $1\{\alpha_n = \alpha_i\}$ is an indicator function that is 1 if $\alpha_n = \alpha_i$, and 0 else. The denominator in (1) is always nonzero since $1\{\alpha_i = \alpha_i\} = 1$ for all $i \in \{1,2,3\}$.

For the scenario illustrated in Fig. 1, player 1 is the only one to see the highly desirable reward of 10 that is currently in the top left location. Player 1 might want to keep this information private to reduce the chance of other players competing for the same location.

B. Prior work on repeated games

Adaptive methods that converge to a *correlated equilibrium* for repeated play of static games are developed in [1][2][3]. Correlated equilibrium in stochastic games is considered in [4][5]. The formulation of the current paper is the most similar to [5], where a game manager helps to achieve correlated and coarse correlated equilibrium in repeated stochastic games. However, the prior work [1][2][3][4][5] does not consider the problem of information sharing, and so the notions of equilibrium they study do not directly apply in the current context. Further, the work on stochastic games in [4][5] considers an ergodic regime, while the current paper considers arbitrary sample paths that are possibly non-ergodic.

C. Inadequacy of standard equilibrium definitions

Standard definitions of Nash equilibrium [6][7], correlated equilibrium [8][9], and coarse correlated equilibrium [10] are inadequate in the scenario of this paper. That is because such equilibrium definitions require the utility of each player i to be at least as large as it would be if player i individually deviated from the intended strategy *while all other players continue to use the intended strategy*. However, if player i deviates by choosing not to share information, the intended strategies of others may no longer be possible because they might rely on this information.

In principle, one could circumvent this difficulty by forcing the repeated game structure to look like a 1-shot game for which standard notions of equilibrium exist. For example, this could be done by allowing each player to make a binary decision at the start that determines whether or not she will share information. Her remaining decisions can be viewed as an element of a strategy space defined over infinite sequences of actions. This approach is taken in [4] by using the concept of *infinitely punishing deviant behavior* (also see discussions in [11]). There, players are assigned strategies that require them to maximally punish any non-conformist by taking actions (for all time) that are solely designed to yield poor utility for the non-conformist. This can lead to an equilibrium because such punishment never occurs (since all players conform out of fear), and so (mathematically) players do not object to having an unused requirement to punish others as part of their decision strategy. However, this approach does not necessarily capture realistic behavior. Human players will *not* spend the rest of their lives punishing a non-conformist. Rather, human players will adapt their behavior to emerging conditions. The mathematical threat of infinite punishment is seen to be a sham that lacks power to realistically influence behavior.

A modified definition of *subgame perfect equilibrium* is often used as an attempt to make punishment threats credible [11]. It requires equilibrium-type conditions to be met even for unused punishment modes of a strategy. Such conditions can often be met by using finite-length punishment modes. However, in many repeated games, such modes can be appended to almost any strategy to endow that strategy with the subgame perfect equilibrium property (see theorems on repeated games with time average utility metrics in [11]).

Overall, while prior notions of equilibrium for repeated games exist and have well defined mathematical properties, they can be complex and difficult for humans to interpret. Arguably, a human player wants a more direct comparison of the suggested strategy with some other reasonable course of action. This paper uses a much simpler "no regret" guarantee that compares the suggestions to a single baseline strategy, rather than to all possible strategies. The baseline strategy is defined by any alternative sequence of actions that a player wants to consider. The challenge is to dynamically make suggestions as the game is being played. The suggestions must meet the desired performance for arbitrary sample paths, and therefore must adapt as new events emerge.

D. The game manager

This paper deviates from the standard equilibrium approach by assuming the existence of a *game manager* to which players pass their information. Further, every round it is assumed that each player $i \in \{1,\ldots,N\}$ has a *baseline decision* $b_i(t) \in \mathcal{A}_i$. The baseline decision $b_i(t)$ is any decision the human player would choose if she were operating on her own without a manager and without information sharing. The manager takes the given information every round t and produces *suggested actions* for each player. The suggested actions must have the property that, if every player uses the suggestions, a concave function of average utilities across players is maximized subject to the constraint that each player receives an average utility at least as large as the average value it would earn if all players used their baseline strategies. The running average

utility gains between the suggested and baseline strategies can be given to each player on each round. The understanding is that players have more incentive to take suggested actions if they see the gains of doing so. This approach does not rely on punishment modes in a complex strategy space.

This formulation is interesting because it defines a specific optimization problem that uses the (possibly human-generated) sequence of baseline strategies as part of the optimization. The solution approach uses Lyapunov optimization theory in this new context. Lyapunov optimization is known to have *universal scheduling* properties that provide analytical guarantees for arbitrary sequences [12]. Those properties are used in the game theory context of this paper to provide a simple online algorithm for making manager decisions as the game is played. This paper shows that the resulting algorithm provides analytical guarantees for arbitrary baseline sequences and arbitrary event sequences, including sequences with no probabilistic description.

II. Examples

To gain intuition, this section provides examples of a 2-player location-reward game with only two locations. The random event vectors $\omega(t) = (\omega_1(t), \omega_2(t))$ are assumed to be independent and identically distributed (i.i.d.) over rounds $t \in \{0, 1, 2, \ldots\}$. This allows exact computation of average utility associated with different strategies. The examples are designed to show that willingness to share information can depend on the statistical distribution of rewards and also on the constraint sets of individual players. For simplicity, the examples in this section assume the full joint probability distribution of $(\omega_1(t), \omega_2(t))$ is known by both players. The general model of Section III treats a more complex scenario where the event vectors $\omega(t)$ are arbitrary sequences, possibly non-ergodic sequences with no known probabilistic structure.

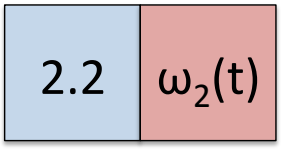

Fig. 2. A location-reward game with two locations and two players. Player 1 knows the reward in the left location, but this reward is always 2.2. Player 2 knows the random reward $\omega_2(t)$ in the right location.

A. Example 1: Beneficially withholding information

Consider a location-reward game (as described in Section I-A) with two players and two locations, as shown in Fig. 2. Player 1 knows the reward $\omega_1(t)$ associated with location 1 and player 2 knows the reward $\omega_2(t)$ associated with location 2. Every round, player 1 can choose from either of the two locations. However, suppose that player 2 is restricted to only choosing location 2. The reward probabilities are:

$$\begin{aligned} \omega_1(t) &= 2.2 \text{ with probability 1} \\ \omega_2(t) &= \begin{cases} 10 & \text{with probability } 1/5 \\ 2 & \text{with probability } 4/5 \end{cases} \end{aligned}$$

The vectors $\omega(t) = (\omega_1(t), \omega_2(t))$ are i.i.d. over rounds t, and the above probabilities are known to both players. Player 2 is the only one with knowledge of the actual realization of the time-varying reward in location 2.

1) Without sharing information: Suppose player 2 does not share its knowledge. If player 1 chooses location 2, its expected utility is $5(1/5)+1(4/5) = 1.8$, which is strictly less than the utility of 2.2 it would achieve by choosing location 1. Hence, the optimal strategy for player 1 is to always choose location 1. Assuming that player 1 uses this optimal strategy, the resulting average utilities are:

$$\overline{u}_1 = 2.2 \quad (2)$$

$$\overline{u}_2 = 10(1/5) + 2(4/5) = 3.6 \quad (3)$$

2) Sharing information: Suppose player 2 chooses to always divulge the value of $\omega_2(t)$ before either player makes a decision. In this case, it is optimal for player 1 to choose location 2 whenever $\omega_2(t) = 10$, and to choose location 1 otherwise. The resulting utilities under this strategy are:

$$\begin{aligned} \overline{u}_1 &= 2.2(4/5) + 5(1/5) = 2.76 \\ \overline{u}_2 &= 2(4/5) + 5(1/5) = 2.6 \end{aligned}$$

In comparison to the utilities of (2)-(3), it is clear that player 1 increases her utility by taking advantage of the shared information. However, player 2 *reduces* her utility because she now competes for the desirable reward of 10 when that reward appears. Thus, in this example, player 2 has no incentive to share information. Player 2 prefers to keep her information private.

B. Example 2: Beneficially sharing information

Consider the same two-player location-reward game as the previous subsection. The only difference is that the reward probabilities are different:

$$\omega_1(t) = 2.2 \text{ with probability 1} \quad (4)$$

$$\omega_2(t) = \begin{cases} 10 & \text{with probability } 1/2 \\ 2 & \text{with probability } 1/2 \end{cases} \quad (5)$$

1) Without sharing information: Without sharing information, it is optimal for player 1 to always choose location 2. The resulting utilities under this strategy are:

$$\overline{u}_1 = \overline{u}_2 = 5(1/2) + 1(1/2) = 3$$

2) Sharing information: If player 2 shares $\omega_2(t)$ with player 1 every round t, then the optimal strategy for player 1 is to choose location 2 whenever $\omega_2(t) = 10$, and choose location 1 otherwise. The resulting utilities under this optimal strategy are:

$$\begin{aligned} \overline{u}_1 &= 2.2(1/2) + 5(1/2) = 3.6 \\ \overline{u}_2 &= 2(1/2) + 5(1/2) = 3.5 \end{aligned}$$

In this example, sharing information allows *both* players to increase their average utility. Player 2 wins by revealing the value of $\omega_2(t)$ because it discourages player 1 from competing when the reward is small.

The examples in Sections II-A and II-B show that changing the probability distribution of random events can impact the optimality or sub-optimality of sharing information. The problem treated in this paper is even more challenging because the probabilities are unknown and can possibly change.

C. Example 3: Unrestricted actions

Now consider the same example of the previous subsection, with the same reward probabilities as (4)-(5). The only difference is that now both players 1 and 2 are free to select either of the two locations.

1) Without sharing information: Without knowledge of $\omega_2(t)$, it is optimal for player 1 to always choose location 2. Indeed, the *smallest* expected reward it can earn by doing this is computed by assuming it must always share its rewards in location 2, and this leads to a value of $5/2+1/2 = 3 > 2.2$. If player 2 assumes that player 1 uses its optimal policy, then its best strategy is to choose location 2 whenever $\omega_2(t) = 10$, and to choose location 1 if $\omega_2(t) = 2$. Under these optimized strategies, the utilities are:

$$\begin{aligned} \overline{u}_1 &= 5(1/2) + 2(1/2) = 3.5 \\ \overline{u}_2 &= 5(1/2) + 2.2(1/2) = 3.6 \end{aligned}$$

2) Sharing information: Suppose player 2 shares $\omega_2(t)$ with player 1 every round t. In this case, it is optimal for both players to choose location 2 whenever $\omega_2(t) = 10$. However, the optimal decisions for players 1 and 2 are unclear when $\omega_2(t) = 2$, since optimality of each player depends on the policy implemented by the other. However, the *highest* utility that player 2 can achieve is 3.6 (which assumes it always gets exclusive access to the 2.2 reward in location 1 when $\omega_2(t) = 2$). Thus, at best, player 2 can only achieve the same utility by sharing its information, but more likely stands to *loose* utility when this information is shared. Therefore, in this example, a rational and self-interested player 2 would not want to share information.

III. The general model

Fix N as an integer larger than 1. Suppose there are N players that play a game over rounds $t \in \{0, 1, 2, \ldots\}$. Define $\mathcal{N} = \{1, \ldots, N\}$ as the set of players. Fix M as a positive integer (possibly different from N) and suppose that $\omega(t) = (\omega_1(t), \omega_2(t), \ldots, \omega_M(t))$ is a sequence of event vectors over rounds $t \in \{0, 1, 2, \ldots\}$. Assume that $\omega(t)$ takes values in some abstract (possibly infinite) set Ω on each round t. The vector sequence $\omega(t)$ is otherwise arbitrary and can have arbitrary correlations over entries and over time. A probabilistic description of $\omega(t)$ is not necessarily known to the players, and such a probabilistic description may not even exist. For each $i \in \mathcal{N}$, define $\mathcal{S}_i$ as the subset of $\{1, \ldots, M\}$ associated with components of $\omega(t)$ that player i can observe at the beginning of each round. That is, player i knows $\omega_j(t)$ for all $j \in \mathcal{S}_i$ and for all t. It is assumed that $\cup_{i=1}^N \mathcal{S}_i = \{1, \ldots, M\}$, so that the combined knowledge of all players gives the full $\omega(t)$ vector.

Every round t, each player $i \in \mathcal{N}$ observes its components of $\omega(t)$ and chooses an *action* $\alpha_i(t)$ as an element in some abstract (possibly infinite) action set $\mathcal{A}_i$. Let $\alpha(t) = (\alpha_1(t), \ldots, \alpha_N(t))$ be the action vector. The resulting utility $u_i(t)$ earned by player i on round t is a general function of $\alpha(t)$ and $\omega(t)$:

$$u_i(t) = \hat{u}_i(\alpha(t), \omega(t)) \ \forall i \in \mathcal{N}$$

A. Assumptions

The functions $\hat{u}_i(\alpha, \omega)$ are assumed to be non-negative and upper-bounded. Specifically, for each $i \in \mathcal{N}$, assume there is a maximum utility value $u_i^{max} < \infty$ such that:

$$0 \leq \hat{u}_i(\alpha, \omega) \leq u_i^{max} \quad \forall (\alpha, \omega) \in \mathcal{A}_1 \times \cdots \times \mathcal{A}_N \times \Omega$$

For simplicity, further assume the utility functions are such that for every $\omega \in \Omega$, the problem of choosing an action vector $\alpha = (\alpha_1, \ldots, \alpha_N)$ to maximize a weighted sum of utilities has a well defined (possibly non-unique) maximizing solution α^*. Specifically, the following problem has a well defined maximum (so that the supremum objective function value is achievable) for all possible real numbers β_i:

$$\text{Maximize:} \quad \textstyle\sum_{i=1}^N \beta_i \hat{u}_i(\alpha, \omega) \tag{6}$$

$$\text{Subject to:} \quad \alpha_i \in \mathcal{A}_i \ \forall i \in \mathcal{N} \tag{7}$$

This is a mild assumption that holds in most practical cases.[1] For example, (6)-(7) is guaranteed to have a maximizing solution $\alpha^* = (\alpha_1^*, \ldots, \alpha_N^*)$ when all sets $\mathcal{A}_i$ are finite. It is also guaranteed to have a well defined maximizer when the sets $\mathcal{A}_i$ are infinite but are compact subsets of a finite-dimensional vector space, and when the utility functions $\hat{u}_i(\alpha, \omega)$ are continuous in α for all $\omega \in \Omega$.

The utility functions are otherwise arbitrary. In particular, they are not required to have convexity or concavity properties.

B. Baseline actions and the game manager

Every round t, each player i observes $\omega_j(t)$ for all $j \in \mathcal{S}_i$ and then makes a *baseline decision* $b_i(t) \in \mathcal{A}_i$. The resulting sequence $\{b_i(t)\}_{t=0}^{\infty}$ can be arbitrary and has no assumed structure. However, the understanding is that the baseline decision on round t is an action that player i would want to take if it did not have access to shared information or to suggestions of a game manager. It can be based on the player i observations of current and past events.

At the beginning of each round t, all players $i \in \mathcal{N}$ privately upload their observations $\omega_j(t)$ (for all $j \in \mathcal{S}_i$) and their baseline decisions $b_i(t)$ to the game manager. Thus, on round t, the manager knows the full $\omega(t)$ vector, all of the baseline decisions $b_i(t)$, and the complete history of past events. It is assumed that the manager has only *causal knowledge*, and hence it does not know the future values $\omega(\tau)$ and $b_i(\tau)$ for $\tau > t$. The manager uses its information on round t to compute *suggested actions* $\tilde{\alpha}_i(t) \in \mathcal{A}_i$ that it delivers to each individual player. Define $\tilde{\alpha}(t) = (\tilde{\alpha}_1(t), \ldots, \tilde{\alpha}_N(t))$ and $b(t) = (b_1(t), \ldots, b_N(t))$.

[1] Assuming existence of a maximizer α^* for the problem (6)-(7) simplifies exposition but is not crucial to the analysis. This assumption can be avoided by using the *C-additive approximation* theory in [12].

C. An (overly?) ambitious optimization problem

Define $u_i(t)$ and $x_i(t)$ for each $i \in \mathcal{N}$ and each $t \in \{0, 1, 2, \ldots\}$ by:

$$\begin{aligned} u_i(t) &= \hat{u}_i(\tilde{\alpha}(t), \omega(t)) \\ x_i(t) &= \hat{u}_i(b(t), \omega(t)) \end{aligned}$$

The value $u_i(t)$ is the utility earned by player i on round t if all players choose the suggested actions, while $x_i(t)$ is the corresponding utility if all players choose their baseline decisions. Define time averages for all $t > 0$ by:

$$\begin{aligned} \overline{u}_i(t) &= \frac{1}{t}\sum_{\tau=0}^{t-1} \hat{u}_i(\tilde{\alpha}(t), \omega(t)) \\ \overline{x}_i(t) &= \frac{1}{t}\sum_{\tau=0}^{t-1} \hat{u}_i(b(t), \omega(t)) \end{aligned}$$

It is useful to introduce an ambitious optimization problem that will be modified later. Subject to the emerging $\omega(t)$ and $b(t)$ sequences, the goal is for the manager to make suggestions $\tilde{\alpha}(t)$ that solve the following:

$$\begin{aligned} \text{Maximize:} \quad & \liminf_{t\to\infty} \phi(\overline{u}_1(t), \ldots, \overline{u}_N(t)) && (8) \\ \text{Subject to:} \quad & \liminf_{t\to\infty} [\overline{u}_i(t) - \overline{x}_i(t)] \geq 0 \ \ \forall i \in \mathcal{N} && (9) \\ & \tilde{\alpha}_i(t) \in \mathcal{A}_i \ \ \forall i \in \mathcal{N}, \forall t \in \{0, 1, 2, \ldots\} && (10) \end{aligned}$$

where $\phi(u_1, \ldots, u_N)$ is a continuous and concave function defined over the hyper-rectangle of all $(u_1, \ldots, u_N) \in \prod[0, u_i^{max}]$. The constraint (9) ensures the time average utility player i receives if all players use the suggestions of the manager is at least as good as the utility it would receive if all players used their baseline strategies.

The *definition* of optimality for (8)-(10) requires a more careful treatment, and this issue is discussed more precisely in the next subsection. In particular, the *causality constraint* of the game manager does not explicitly appear anywhere in (8)-(10). Regardless, it can be shown that the above problem is always *feasible*, so that it is always possible to satisfy constraints (9)-(10) via a simple causal algorithm for game manager decisions: Consider the trivial decisions $(\tilde{\alpha}_1(t), \ldots, \tilde{\alpha}_N(t)) = (b_1(t), \ldots, b_N(t))$ for all $t \in \{0, 1, 2, \ldots\}$. This means that the game manager suggests nothing more than the baseline actions for all rounds t. These suggestions are implementable in a causal manner because they only use the baseline decisions given to the game manager at the beginning of each round. Since $b_i(t) \in \mathcal{A}_i$ for all i and all t, one has $\tilde{\alpha}_i(t) = b_i(t) \in \mathcal{A}_i$ for all i, and so the constraints (10) are satisfied. Further, this trivial suggestion strategy immediately implies $u_i(t) = x_i(t)$ for all t, and so constraint (9) is trivially satisfied.

Thus, it is *always possible* for a game manager to achieve the constraints (9)-(10). The constraints (9) provide rational players an incentive to use the suggested actions, since it ensures the resulting time average utilities are at least as good as those of the baseline decisions. Thus, it is assumed throughout this paper that all players choose the suggestions of the manager, so that $\tilde{\alpha}_i(t) = \alpha_i(t)$ for all i and all t. In particular, one has for all rounds t:

$$\begin{aligned} u_i(t) &= \hat{u}_i(\alpha(t), \omega(t)) && (11) \\ x_i(t) &= \hat{u}_i(b(t), \omega(t)) && (12) \end{aligned}$$

Define:

$$\begin{aligned} u(t) &= (u_1(t), \ldots, u_N(t)) \\ \overline{u}(t) &= (\overline{u}_1(t), \ldots, \overline{u}_N(t)) \end{aligned}$$

Under the assumption $\tilde{\alpha}_i(t) = \alpha_i(t)$ for all i and all t, the problem becomes:

$$\begin{aligned} \text{Maximize:} \quad & \liminf_{t\to\infty} \phi(\overline{u}(t)) && (13) \\ \text{Subject to:} \quad & \liminf_{t\to\infty} [\overline{u}_i(t) - \overline{x}_i(t)] \geq 0 \ \ \forall i \in \mathcal{N} && (14) \\ & \alpha_i(t) \in \mathcal{A}_i \ \ \forall i \in \mathcal{N}, \forall t \in \{0, 1, 2, \ldots\} && (15) \end{aligned}$$

where $u_i(t)$ and $x_i(t)$ are defined in (11) and (12). The problem (13)-(15) shall be referred to as the *infinite future knowledge optimization problem*.

D. A modified objective

The infinite sequences $\omega(t)$ and $b(t)$ for $t \in \{0, 1, 2, \ldots\}$ are arbitrary. In practice, these sequences might depend on previous suggestions of the game manager. However, the problem (13)-(15) does not specify how future values of $\omega(\tau)$ and $b(\tau)$ depend on control decisions. Thus, to understand optimal utility in (13)-(15), it useful to view $\{\omega(t)\}_{t=0}^{\infty}$ and $\{b(t)\}_{t=0}^{\infty}$ as arbitrary sequences that are defined at the start of the game but with values that are only sequentially revealed as the game progresses. In this way, future values of $\omega(\tau)$ and $b(\tau)$ are not influenced by the emerging decisions of the manager.

With this structure, optimality of (13)-(15) is defined over all possible action sequences $\{\alpha(t)\}_{t=0}^{\infty}$ that provide utilities with respect to the given $\{\omega(t)\}_{t=0}^{\infty}$ and $\{b(t)\}_{t=0}^{\infty}$ sequences. In principle, the supremum time average utility in (13) can be computed offline based on non-causal knowledge of the full sequences $\{\omega(t)\}_{t=0}^{\infty}$ and $\{b(t)\}_{t=0}^{\infty}$. Since optimality is defined in terms of full knowledge of the future, the problem (13)-(15) is called the *infinite future knowledge optimization problem*. It is not clear if the supremum objective function value for this problem can be achieved by a practical algorithm that makes causal decisions. The stochastic optimization theory in [12] shows that the supremum *can* be achieved in a causal manner in the special case when the sequences $\{\omega(t)\}_{t=0}^{\infty}$ and $\{b(t)\}_{t=0}^{\infty}$ have an *ergodic* structure.[2] That is because, in the ergodic case, the problem (13)-(15) fits into the general framework of [12].

However, this paper considers problems with possibly non-ergodic input sequences $\{\omega(t)\}_{t=0}^{\infty}$ and $\{b(t)\}_{t=0}^{\infty}$. In this context, it is not clear if the supremum in (13)-(15) can be causally achieved. Remarkably, this paper shows that a modified objective, defined in terms of a finite but arbitrarily large window of time in which the future is known, *can in fact be causally achieved* (to within any arbitrarily small but

[2]Specifically, [12] shows optimality can be achieved (arbitrarily closely) in the case when the random event process is modulated by a finite state irreducible (possibly periodic) discrete time Markov chain.

positive error). This is done using the *T-slot lookahead utility* developed in [12]: Fix T as a positive integer and partition the rounds $t \in \{0, 1, 2, \ldots\}$ into successive frames of size T, so that frame k consists of rounds $\{kT, \ldots, (k+1)T - 1\}$ (for each $k \in \{0, 1, 2, \ldots\}$). For each round k and for given realizations of $\{\omega(t)\}_{t=0}^{\infty}$ and $\{b(t)\}_{t=0}^{\infty}$, define $\psi_T[k]$ as the supremum objective value in the following optimization problem, optimized over all choices of the decision vector $\alpha(t) = (\alpha_1(t), \ldots, \alpha_N(t))$:

$$\text{Max:} \quad \phi(\gamma_1, \ldots, \gamma_N) \tag{16}$$

$$\text{Subj to:} \quad \gamma_i = \frac{1}{T}\sum_{\tau=kT}^{(k+1)T-1} u_i(\alpha(\tau), \omega(\tau)) \ \ \forall i \in \mathcal{N} \tag{17}$$

$$\gamma_i \geq \frac{1}{T}\sum_{\tau=kT}^{(k+1)T-1} u_i(b(\tau), \omega(\tau)) \ \ \forall i \in \mathcal{N} \tag{18}$$

$$\alpha_i(\tau) \in \mathcal{A}_i \ \ \forall \tau \in \{kT, \ldots, (k+1)T - 1\} \tag{19}$$

$$\gamma_i \in [0, u_i^{max}] \ \ \forall i \in \mathcal{N} \tag{20}$$

In particular, if $(\gamma_1^*, \ldots, \gamma_N^*)$ and $\{\alpha^*(\tau)\}_{\tau=kT}^{(k+1)T-1}$ forms an optimal solution to the above problem, then $\psi_T[k] = \phi(\gamma_1^*, \ldots, \gamma_N^*)$. The above problem is always feasible with a well defined and finite supremum $\psi_T[k]$. In general, the supremum may not be achievable. However, for all $\epsilon > 0$, there are vectors $(\gamma_1^*, \ldots, \gamma_N^*)$ and $\{\alpha^*(\tau)\}_{\tau=kT}^{(k+1)T-1}$ that satisfy the constraints of the above problem and that satisfy:

$$\psi_T[k] - \epsilon \leq \phi(\gamma_1^*, \ldots, \gamma_N^*) \leq \psi_T[k]$$

The value $\psi_T[k]$ represents the maximum average utility achievable over frame k provided that the regret constraints are satisfied by averages over that frame, and assuming that future values of the vectors $b(t)$ and $\omega(t)$ are fully known over the frame. Since $\psi_T[k]$ requires knowledge of the future to compute, it seems unlikely that a practical algorithm could achieve performance that is competitive with the $\psi_T[k]$ values. Remarkably, this paper develops a *causal* algorithm (without requiring knowledge of the future) that satisfies the constraints (14)-(15) and that achieves:

$$\liminf_{K\to\infty}\left[\overline{u}_i(KT) - \frac{1}{K}\sum_{k=0}^{K-1}\psi_T[k]\right] \geq -BT/V$$

where B is a constant and V is a parameter that can be chosen as large as desired (with a tradeoff in the corresponding convergence time). This holds under the same algorithm *for all possible (finite) values of* T. In particular, for any desired value of T, it is possible to choose a sufficiently large value of V so that long term performance is arbitrarily close to the average of the $\psi_T[k]$ values.

It can be shown that, for any T and for any $\{\omega(t)\}_{t=0}^{\infty}$ and $\{b(t)\}_{t=0}^{\infty}$ sequences, the $\liminf$ average of $\psi_T[k]$ is less than or equal the optimal objective function value in (13) for the infinite future knowledge optimization problem. In cases when sequences $\{\omega(t)\}_{t=0}^{\infty}$ and $\{b(t)\}_{t=0}^{\infty}$ are ergodic and mild additional assumptions are satisfied, the limiting average of $\psi_T[k]$ converges to this value as $T \to \infty$. However, this is not true in general non-ergodic situations. In particular, example sequences can be given for which the optimal value in (13) is strictly larger than the following value:

$$\liminf_{T\to\infty}\left[\liminf_{K\to\infty}\frac{1}{K}\sum_{k=0}^{K-1}\psi_T[k]\right]$$

Nevertheless, the average of $\psi_T[k]$ still provides a meaningful and challenging utility target.

IV. Weighted sum of utilities

This subsection considers a weighted sum of utilities, so that:

$$\phi(u_1, \ldots, u_N) = \sum_{i=1}^{N}\theta_i u_i$$

for given real numbers θ_i. Subsection V considers the more general case when $\phi(\gamma_1, \ldots, \gamma_N)$ is any concave function.

A. Virtual queues

To achieve the time average constraint (14), for each $i \in \mathcal{N}$ define a *virtual queue* $Q_i(t)$ that is initialized to $Q_i(0) = 0$ and that has update equation:

$$Q_i(t+1) = \max[Q_i(t) + x_i(t) - u_i(t), 0] \tag{21}$$

where $u_i(t) = \hat{u}_i(\alpha(t), \omega(t))$ and $x_i(t) = \hat{u}(\alpha(t), \omega(t))$. These virtual queues are kept in the game manager and updated at the end of every round t based on its suggestion $\alpha(t)$ and on knowledge of the $\omega(t)$ vector. Recall that the manager knows all entries of $\omega(t)$ since each player i tells it the values of $\omega_j(t)$ for $j \in \mathcal{S}_i$, and $\cup_{i\in\mathcal{N}}\mathcal{S}_i = \{1, \ldots, M\}$.

The following queueing lemma is standard for stochastic network optimization [12]. The proof is given for completeness.

Lemma 1: (Virtual queues [12]) If $Q_i(t)$ satisfies (21) then:
a) For all $t \in \{1, 2, 3, \ldots\}$ one has:

$$\overline{u}_i(t) \geq \overline{x}_i(t) - \left[\frac{Q_i(t) - Q_i(0)}{t}\right]$$

b) If $\lim_{t\to\infty} Q_i(t)/t = 0$ for all $i \in \mathcal{N}$, then constraints (14) are satisfied.

Proof: From (21) one has for all $\tau \in \{0, 1, 2, \ldots\}$:

$$Q_i(\tau+1) \geq Q_i(\tau) + x_i(\tau) - u_i(\tau)$$

Thus:

$$Q_i(\tau+1) - Q_i(\tau) \geq x_i(\tau) - u_i(\tau)$$

Fix $t > 0$. Summing the above over $\tau \in \{0, 1, \ldots, t-1\}$ gives:

$$Q_i(t) - Q_i(0) \geq \sum_{\tau=0}^{t-1}[x_i(\tau) - u_i(\tau)]$$

Dividing the result by t and rearranging terms gives the result of part (a). Part (b) immediately follows. □

A queue $Q_i(t)$ that satisfies $\lim_{t\to\infty} Q_i(t)/t = 0$ with probability 1 is said to be *rate stable*.

B. Drift-plus-penalty

Define $Q(t) = (Q_1(t), \ldots, Q_N(t))$ and define $||Q(t)||^2 = \sum_{i\in\mathcal{N}} Q_i(t)^2$. Define $L(t) = \frac{1}{2}||Q(t)||^2$, called a *Lyapunov function*. Define $\Delta(t) = L(t+1) - L(t)$. The *drift-plus-penalty* method of [12] observes $Q(t)$, $\omega(t)$, and $b(t)$ every round t, and then takes a control action $\alpha(t)$ to minimize a bound on the following expression:

$$\Delta(t) - V\sum_{i=1}^{N}\theta_i u_i(t)$$

where V is a non-negative parameter that affects a performance tradeoff. The intuition is as follows: Including the drift term $\Delta(t)$ in the above minimization maintains stable queues so that the time average constraints can be satisfied. Including the penalty term $-V\sum_{i=1}^{N}\theta_i u_i(t)$ encourages the controller to make decisions that give a desirable weighted sum of utilities. Using larger values of V places more emphasis on this "penalty minimization."

Lemma 2: Under any algorithm for choosing $\alpha(t) \in \mathcal{A}_1 \times \cdots \times \mathcal{A}_N$, one has for all $t \in \{0, 1, 2, \ldots\}$:

$$\begin{aligned}&\Delta(t) - V\sum_{i=1}^{N}\theta_i u_i(t)\\ &\leq B - V\sum_{i=1}^{N}\theta_i\hat{u}_i(\alpha(t),\omega(t))\\ &+\sum_{i=1}^{N}Q_i(t)\left[\hat{u}_i(b(t),\omega(t)) - \hat{u}_i(\alpha(t),\omega(t))\right] \qquad (22)\end{aligned}$$

where $B = \frac{1}{2}\sum_{i=1}^{N}(u_i^{max})^2$.

Proof: Squaring (21) and using $\max[y, 0]^2 \leq y^2$ gives:

$$Q_i(t+1)^2 \leq Q_i(t)^2 + (x_i(t) - u_i(t))^2 + 2Q_i(t)[x_i(t) - u_i(t)]$$

Summing over $i \in \{1, \ldots, N\}$ and dividing by 2 gives:

$$\begin{aligned}\Delta(t) &\leq \frac{1}{2}\sum_{i=1}^{N}(x_i(t) - u_i(t))^2 + \sum_{i=1}^{N}Q_i(t)[x_i(t) - u_i(t)]\\ &\leq B + \sum_{i=1}^{N}Q_i(t)[x_i(t) - u_i(t)]\end{aligned}$$

where the last inequality follows because $(x_i(t) - u_i(t)) \in [-u_i^{max}, u_i^{max}]$ for all t. Subtracting $V\sum_{i=1}^{N}\theta_i u_i(t)$ from both sides and using $u_i(t) = \hat{u}_i(\alpha(t),\omega(t))$ and $x_i(t) = \hat{u}_i(b(t),\omega(t))$ gives the result. □

The algorithm takes actions every round t to greedily minimize the right-hand-side of the drift-plus-penalty expression (22). The only terms on the right-hand-side of (22) that are affected by control decisions $\alpha(t)$ are:

$$-\sum_{i=1}^{N}V\theta_i\hat{u}_i(\alpha(t),\omega(t)) - \sum_{i=1}^{N}Q_i(t)\hat{u}_i(\alpha(t),\omega(t))$$

The resulting algorithm is as follows: Initialize $Q_i(0) = 0$ for $i \in \mathcal{N}$. Every round $t \in \{0, 1, 2, \ldots\}$, the game manager observes vectors $Q(t)$, $\omega(t)$, and $b(t)$ and does the following:

- (Decisions) Choose suggestion vector $\alpha(t) = (\alpha_1(t), \ldots, \alpha_N(t))$ as the solution to the following:

  Maximize: $\sum_{i=1}^{N}\hat{u}_i(\alpha(t),\omega(t))[V\theta_i + Q_i(t)]$
  Subject to: $\alpha_i(t) \in \mathcal{A}_i \ \forall i \in \mathcal{N}$
- (Send suggestions) Send suggestions $\alpha_i(t)$ to each player $i \in \mathcal{N}$.
- (Queue update) For each $i \in \mathcal{N}$, update $Q_i(t)$ via (21).

It is clear that the above algorithm is causal: It only requires knowledge of the current $Q(t), \omega(t), b(t)$ and does not require knowledge of the future.

C. Performance for weighted utilities

The following theorem shows that the algorithm satisfies constraints (14)-(15) with constraint violations that decay like $O(\sqrt{V/t})$.

Theorem 1: (Constraint satisfaction) Fix $V \geq 0$ and assume the algorithm in the previous subsection is used with this V and with $Q_i(0) = 0$ for all $i \in \mathcal{N}$. For all $t \in \{1, 2, 3, \ldots\}$ one has:

a) $\frac{||Q(t)||}{t} \leq \sqrt{\frac{2B+2V\sum_{i=1}^{N}|\theta_i|u_i^{max}}{t}}$.

b) $\overline{u}_i(t) - \overline{x}_i(t) \geq -\sqrt{\frac{2B+2V\sum_{i=1}^{N}|\theta_i|u_i^{max}}{t}}$.

c) The constraints (14)-(15) are satisfied.

Proof: Fix $\tau \in \{0, 1, 2, \ldots\}$. Since the algorithm makes decisions for $\alpha(\tau)$ to minimize the right-hand-side of (22), one has:

$$\begin{aligned}&\Delta(\tau) - V\sum_{i=1}^{N}\theta_i u_i(\tau)\\ &\leq B - V\sum_{i=1}^{N}\theta_i\hat{u}_i(\alpha^*(\tau),\omega(\tau))\\ &+\sum_{i=1}^{N}Q_i(\tau)[\hat{u}_i(b(\tau),\omega(\tau)) - \hat{u}_i(\alpha^*(\tau),\omega(\tau))] \qquad (23)\end{aligned}$$

where $\alpha^*(\tau) = (\alpha_1^*(\tau), \ldots, \alpha_N^*(\tau))$ is any alternative vector that satisfies $\alpha_i^*(\tau) \in \mathcal{A}_i$ for all $i \in \mathcal{N}$. A valid choice is $\alpha^*(\tau) = b(\tau)$. Substituting $\alpha^*(\tau) = b(\tau)$ into the right-hand-side of (23) gives:

$$\Delta(\tau) - V\sum_{i=1}^{N}\theta_i u_i(\tau) \leq B - V\sum_{i=1}^{N}\theta_i\hat{u}_i(b(\tau),\omega(\tau))$$

Rearranging terms gives:

$$\Delta(\tau) \leq B + VC \qquad (24)$$

where $C = \sum_{i=1}^{N}|\theta_i|u_i^{max}$. The above holds for all rounds $\tau \in \{0, 1, 2, \ldots\}$. Fix $t > 0$. Summing (24) over $\tau \in \{0, \ldots, t-1\}$ gives:

$$L(t) - L(0) \leq (B + VC)t$$

Since $L(t) = \frac{1}{2}||Q(t)||^2$ and $L(0) = 0$ one has:

$$||Q(t)||^2 \leq 2(B + VC)t$$

Dividing by t^2 and taking square roots gives:

$$\frac{||Q(t)||}{t} \leq \sqrt{\frac{2(B+VC)}{t}}$$

This proves part (a). Part (b) follows from (a) together with Lemma 1. Part (c) follows directly from (b). □

The above theorem does not use a value T. The value T is also never used in the algorithm implementation. It is only used in the performance theorem below.

Theorem 2: (Performance) Fix $V \geq 0$ and assume the algorithm in the previous subsection is used with this V and with $Q_i(0) = 0$ for all $i \in \mathcal{N}$. For all positive integers T and K the following holds:

$$\sum_{i=1}^{N} \theta_i \overline{u}_i(KT) \geq \frac{1}{K} \sum_{k=0}^{K-1} \psi_T[k] - \frac{TB}{V}$$

Theorem 2 shows that for any frame size T and for any number of frames K, the weighted sum of utilities achieved by this algorithm over the first T frames is at most TB/V less than the average of the ideal T-slot lookahead values $\psi_T[k]$ over those frames. Notice that this holds for all frame sizes T. Since the algorithm does not use a value T as input, the above theorem can be viewed as a class of performance bounds that are parameterized by T, all of which are satisfied. The error term TB/V can be made as close to 0 as desired by choosing V appropriately large. This is remarkable, particularly when T is large, because the ideal $\psi_T[k]$ value is defined in terms of perfect knowledge of the future over T rounds, whereas the algorithm does not know the future. The tradeoff is that a large value of V affects the *convergence time* required to meet the desired constraints, as specified by part (b) of Theorem 1.

Proof: (Theorem 2) Fix k as a non-negative integer. Summing (23) over $\tau \in \{kT, \ldots, (k+1)T-1\}$ gives:

$$\begin{aligned}
&\sum_{\tau=kT}^{(k+1)T-1} \Delta(\tau) - V \sum_{\tau=kT}^{(k+1)T-1} \sum_{i=1}^{N} \theta_i u_i(\tau) \\
&\leq TB - V \sum_{\tau=kT}^{(k+1)T-1} \sum_{i=1}^{N} \theta_i \hat{u}_i(\alpha^*(\tau), \omega(\tau)) \\
&+ \sum_{\tau=kT}^{(k+1)T-1} \sum_{i=1}^{N} Q_i(\tau)[\hat{u}_i(b(\tau), \omega(\tau)) - \hat{u}_i(\alpha^*(\tau), \omega(\tau))] \\
&\leq TB + 2B \sum_{m=0}^{T-1} m - V \sum_{\tau=kT}^{(k+1)T-1} \sum_{i=1}^{N} \theta_i \hat{u}_i(\alpha^*(\tau), \omega(\tau)) \\
&+ \sum_{\tau=kT}^{(k+1)T-1} \sum_{i=1}^{N} Q_i(kT)[\hat{u}_i(b(\tau), \omega(\tau)) - \hat{u}_i(\alpha^*(\tau), \omega(\tau))] \\
&= T^2 B - VT \sum_{i=1}^{N} \theta_i \left[\frac{1}{T} \sum_{\tau=kT}^{(k+1)T-1} \hat{u}_i(\alpha^*(\tau), \omega(\tau)) \right] \\
&+ \sum_{\tau=kT}^{(k+1)T-1} \sum_{i=1}^{N} Q_i(kT)[\hat{u}_i(b(\tau), \omega(\tau)) - \hat{u}_i(\alpha^*(\tau), \omega(\tau))]
\end{aligned}$$

Now define $\gamma^* = (\gamma_1^*, \ldots, \gamma_N^*)$ by:

$$\gamma_i^* = \frac{1}{T} \sum_{\tau=kT}^{(k+1)T-1} \hat{u}_i(\alpha^*(\tau), \omega(\tau)) \ \forall i \in \mathcal{N}$$

Then:

$$\begin{aligned}
&\sum_{\tau=kT}^{(k+1)T-1} \Delta(\tau) - V \sum_{\tau=kT}^{(k+1)T-1} \sum_{i=1}^{N} \theta_i u_i(\tau) \\
&\leq T^2 B - VT \sum_{i=1}^{N} \theta_i \gamma_i^* \\
&+ T \sum_{i=1}^{N} Q_k(kT) \left[\frac{1}{T} \sum_{\tau=kT}^{(k+1)T-1} \hat{u}_i(b(\tau), \omega(\tau)) - \gamma_i^* \right]
\end{aligned}$$

Now fix $\epsilon > 0$ and define $\alpha^*(\tau)$ for $\tau \in \{kT, \ldots, (k+1)T-1\}$ as decisions that satisfy the constraints (17)-(20) and yield:

$$\psi_T[k] - \epsilon \leq \sum_{i=1}^{N} \theta_i \gamma_i^* \leq \psi_T[k]$$

It follows that:

$$\begin{aligned}
&\sum_{\tau=kT}^{(k+1)T-1} \Delta(\tau) - V \sum_{\tau=kT}^{(k+1)T-1} \sum_{i=1}^{N} \theta_i u_i(\tau) \qquad (25) \\
&\leq T^2 B - VT\psi_T[k] + \epsilon TV \qquad (26)
\end{aligned}$$

This holds for all $\epsilon > 0$. Taking a limit as $\epsilon \to 0$ gives:

$$\begin{aligned}
&\sum_{\tau=kT}^{(k+1)T-1} \Delta(\tau) - V \sum_{\tau=kT}^{(k+1)T-1} \sum_{i=1}^{N} \theta_i u_i(\tau) \\
&\leq T^2 B - VT\psi_T[k]
\end{aligned}$$

Fix a positive integer K. Summing the above over $k \in \{0, 1, \ldots, K-1\}$ gives:

$$\begin{aligned}
&\sum_{\tau=0}^{KT-1} \Delta(\tau) - V \sum_{\tau=0}^{KT-1} \sum_{i=1}^{N} \theta_i u_i(\tau) \\
&\leq T^2 BK - VT \sum_{k=0}^{K-1} \psi_T[k]
\end{aligned}$$

That is:

$$\begin{aligned}
&L(KT) - L(0) - VKT \sum_{i=1}^{N} \theta_i \overline{u}_i(KT) \\
&\leq T^2 BK - VKT \frac{1}{K} \sum_{k=0}^{K-1} \psi_T[k]
\end{aligned}$$

Dividing by VKT and using the fact that $L(0) = 0$ and $L(KT) \geq 0$ gives:

$$-\sum_{i=1}^{N} \theta_i \overline{u}_i(KT) \leq \frac{TB}{V} - \frac{1}{K} \sum_{k=0}^{K-1} \psi_T[k]$$

Therefore:

$$\sum_{i=1}^{N} \theta_i \overline{u}_i(KT) \geq \frac{1}{K} \sum_{k=0}^{K-1} \psi_T[k] - \frac{TB}{V}$$

□

D. Discussion

Theorems 1 and 2 are *deterministic results* that are guaranteed hold on every sample path, regardless of the probability model. Theorem 2 holds for arbitrarily large values of T and shows that, for large K, averages over KT rounds give time average performance that is arbitrarily close to the value $\frac{1}{K}\sum_{k=0}^{K-1}\psi_T[k]$, where $\psi_T[k]$ is an ideal value based on knowledge of T rounds into the future. If the $\{b(t)\}_{t=0}^{\infty}$ and $\{\omega(t)\}_{t=0}^{\infty}$ processes are ergodic and mild additional assumptions are satisfied, then the value $\frac{1}{K}\sum_{k=0}^{K-1}\psi_T[k]$ approaches the ergodic supremum of the problem (13)-(15) as $T\to\infty$ and $K\to\infty$. Intuitively, this is why the algorithm comes arbitrarily close to the solution of (13)-(15) in the ergodic case.

However, the solution of (13)-(15) requires perfect knowledge of *all time into the future*, rather than just a finite horizon of T slots into the future. Thus, in the general non-ergodic case, the value $\frac{1}{K}\sum_{k=0}^{K-1}\psi_T[k]$ does not necessarily come close to the optimal objective function value in the problem (13)-(15), regardless of how large the T value is chosen to be. Nevertheless, the value $\frac{1}{K}\sum_{k=0}^{K-1}\psi_T[k]$ is an intuitively ambitious target, and it is remarkable that, in all cases, this target can be achieved (arbitrarily closely) by a causal algorithm.

V. Concave function of utilities

The previous section considered a weighted sum of utilities. This section considers a general concave function of utilities. Specifically, it uses a function $\phi(u_1,\ldots,u_N)$ that is concave and continuous over $(u_1,\ldots,u_N)\in\prod_{i=1}^{N}[0,u_i^{max}]$. Thus, it treats the general problem (13)-(15). For simplicity of exposition, assume the function $\phi(\cdot)$ is non-negative over $\prod_{i=1}^{N}[0,u_i^{max}]$ (else, a positive constant can be added to it to make it non-negative).

The idea is to introduce *proxy variables* $\gamma_i(t)$ that relate to running averages of player i utility. The proxy variables are chosen every round t in the interval $[0,u_i^{max}]$ and must satisfy:

$$\lim_{t\to\infty}[\overline{\gamma}_i(t)-\overline{u}_i(t)]=0\ \forall i\in\mathcal{N}$$

The above constraint is enforced by defining virtual queues $Z_i(t)$ for all $i\in\mathcal{N}$ with update equation:

$$Z_i(t+1)=Z_i(t)+\gamma_i(t)-u_i(t) \tag{27}$$

By summing (27), is clear that for all $t\in\{1,2,3,\ldots\}$ one has:

$$\overline{u}_i(t)=\overline{\gamma}_i(t)-\frac{Z_i(t)-Z_i(0)}{t} \tag{28}$$

Therefore, it is desirable to make each queue $Z_i(t)$ rate stable. The equality (28) implies:

$$||\overline{u}(t)-\overline{\gamma}(t)||=\frac{||Z(t)-Z(0)||}{t} \tag{29}$$

where:

$$\begin{aligned}\overline{u}(t) &= (\overline{u}_1(t),\ldots,\overline{u}_N(t))\\ \overline{\gamma}(t) &= (\overline{\gamma}_1(t),\ldots,\overline{\gamma}_N(t))\end{aligned}$$

Let $Z(t)=(Z_1(t),\ldots,Z_N(t))$ be the vector of $Z_i(t)$ values. As before, define $Q(t)=(Q_1(t),\ldots,Q_N(t))$, with queues $Q_i(t)$ defined in (21). Define a new Lyapunov function that considers both types of queues:

$$L(t)=\frac{1}{2}||Q(t)||^2+\frac{1}{2}||Z(t)||^2$$

Define $\Delta(t)=L(t+1)-L(t)$. The first step is to compute a bound on the new drift-plus-penalty expression:

$$\Delta(t)-V\phi(\gamma(t))$$

where $\gamma(t)=(\gamma_1(t),\ldots,\gamma_N(t))$.

Lemma 3: Under any algorithm for choosing $\alpha(t)\in\mathcal{A}_1\times\cdots\mathcal{A}_N$ and $\gamma(t)\in\prod_{i=1}^{N}[0,u_i^{max}]$, one has for all $t\in\{0,1,2,\ldots\}$:

$$\begin{aligned}\Delta(t)-V\phi(\gamma(t)) &\le C-V\phi(\gamma(t))\\ &+\sum_{i=1}^{N}Q_i(t)[\hat{u}_i(b(t),\omega(t))-\hat{u}_i(\alpha(t),\omega(t))]\\ &+\sum_{i=1}^{N}Z_i(t)[\gamma_i(t)-\hat{u}_i(\alpha(t),\omega(t))]\end{aligned} \tag{30}$$

where $C=\sum_{i=1}^{N}(u_i^{max})^2$.

Proof: The proof is similar to that of Lemma 2 and is omitted for brevity. □

A. General algorithm

The algorithm makes greedy decisions to minimize the right-hand-side of (30) on every round t. Specifically, every round $t\in\{0,1,2,\ldots\}$, the game manager observes $Q(t)$, $Z(t)$, $\omega(t)$, and $b(t)$ and does the following:

- (Proxy variables) Choose $\gamma(t)=(\gamma_1(t),\ldots,\gamma_N(t))$ as the solution to:

$$\begin{aligned}\text{Maximize:}\quad & V\phi(\gamma(t))-\textstyle\sum_{i=1}^{N}Z_i(t)\gamma_i(t)\\ \text{Subject to:}\quad & 0\le\gamma_i(t)\le u_i^{max}\ \forall i\in\mathcal{N}\end{aligned}$$

- (Suggestions) Choose $\alpha(t)=(\alpha_1(t),\ldots,\alpha_N(t))$ as the solution to:

$$\begin{aligned}\text{Maximize:}\quad & \textstyle\sum_{i=1}^{N}\hat{u}_i(\alpha(t),\omega(t))[Q_i(t)+Z_i(t)]\\ \text{Subject to:}\quad & \alpha_i(t)\in\mathcal{A}_i\ \forall i\in\mathcal{N}\end{aligned}$$

  Next, for each $i\in\mathcal{N}$, send suggestion $\alpha_i(t)$ to player i.
- (Queue update) For each $i\in\mathcal{N}$, update $Q_i(t)$ and $Z_i(t)$ via (21) and (27).

This is again a simple causal algorithm that is implemented as the game progresses.

B. Constraint analysis

Define ϕ^{max} as the maximum of $\phi(\gamma)$ over $\gamma\in\prod_{i=1}^{N}[0,u_i^{max}]$. Such a maximum exists since it is defined over a compact subset of $\mathbb{R}^N$ and the function $\phi(\gamma)$ is continuous. It is known that every continuous function over a compact set is *Lipshitz continuous*, so that there is a positive value M such that:

$$|\phi(\gamma)-\phi(r)|\le M||\gamma-r|| \tag{31}$$

for all $\gamma, r \in \prod_{i=1}^{N}[0, u_i^{max}]$.

Theorem 3: Fix $V \geq 0$ and assume the algorithm in the previous subsection is used with this V and with $Q_i(0) = Z_i(0) = 0$ for all $i \in \mathcal{N}$. For all $t \in \{1, 2, 3, \ldots\}$ one has:

a) $\frac{\sqrt{\sum_{i=1}^{N} Q_i(t)^2 + Z_i(t)^2}}{t} \leq \sqrt{\frac{2C + 2V\phi^{max}}{t}}$

b) $\overline{u}_i(t) - \overline{x}_i(t) \geq -\sqrt{\frac{2C+2V\phi^{max}}{t}}$.

c) The constraints (14)-(15) are satisfied.

d) The utilities satisfy:

$$\phi(\overline{u}(t)) \geq \frac{1}{t}\sum_{\tau=0}^{t-1}\phi(\gamma(\tau)) - M\sqrt{\frac{2C + 2V\phi^{max}}{t}}$$

where $\overline{u}(t) = (\overline{u}_1(t), \ldots, \overline{u}_N(t))$ and M is the Lipschitz constant in (31).

Proof: Since the algorithm makes decisions $\gamma(t)$ and $\alpha(t)$ to minimize the right-hand-side of (30) every round, one has for all rounds $\tau \in \{0, 1, 2, \ldots\}$:

$$\begin{aligned}\Delta(\tau) - V\phi(\gamma(\tau)) &\leq C - V\phi(\gamma^*) \\ &+ \sum_{i=1}^{N} Q_i(\tau)[\hat{u}_i(b(\tau), \omega(\tau)) - \hat{u}_i(\alpha^*(\tau), \omega(\tau))] \\ &+ \sum_{i=1}^{N} Z_i(\tau)[\gamma_i^* - \hat{u}_i(\alpha^*(\tau), \omega(\tau))]\end{aligned}$$

for any alternative vectors $\gamma^* \in \prod_{i=1}^{N}[0, u_i^{max}]$ and $\alpha^*(\tau) \in \mathcal{A}_1 \times \ldots \times \mathcal{A}_N$. Define $\alpha^*(\tau) = b(\tau)$ and $\gamma^* = (\gamma_1^*, \ldots, \gamma_N^*)$ where $\gamma_i^* = \hat{u}_i(\alpha^*(\tau), \omega(\tau))$. Then the above inequality becomes:

$$\Delta(\tau) - V\phi(\gamma(\tau)) \leq C - V\phi(\gamma^*)$$

Rearranging terms and using the fact that $0 \leq \phi(\gamma) \leq \phi^{max}$ for all $\gamma \in \prod_{i=1}^{N}[0, u_i^{max}]$ gives the following for all rounds τ:

$$\Delta(\tau) \leq C + V\phi^{max}$$

The result of part (a) then follows by an argument similar to that of Theorem 1.

Part (b) follows from (a) together with (28). Part (c) follows immediately from part (b) by taking a limit. To prove part (d), note by Jensen's inequality for concave functions that for any round $t > 0$:

$$\begin{aligned}\frac{1}{t}\sum_{\tau=0}^{t-1}\phi(\gamma(\tau)) &\leq \phi(\overline{\gamma}(t)) \\ &\leq \phi(\overline{u}(t)) + M||\overline{\gamma}(t) - \overline{u}(t)|| \\ &\leq \phi(\overline{u}(t)) + M\sqrt{\frac{2C + 2V\phi^{max}}{t}}\end{aligned}$$

where the second inequality above follows by (31) and the final inequality follows by part (a) together with (29). □

Theorem 3 provides a bound on the achieved performance $\phi(\overline{u}(t))$ in terms of the time average $\frac{1}{t}\sum_{\tau=0}^{t-1}\phi(\gamma(\tau))$. The next theorem completes the analysis by bounding the performance of this time average in terms of averages of the ideal $\psi_T[k]$ values over T-slot frames.

Theorem 4: Fix $V \geq 0$ and assume the algorithm in the previous subsection is used with this V and with $Q_i(0) = Z_i(0) = 0$ for all $i \in \mathcal{N}$. For all positive integers T and K the following holds:

$$\frac{1}{KT}\sum_{\tau=0}^{KT-1}\phi(\gamma(\tau)) \geq \frac{1}{K}\sum_{k=0}^{K-1}\psi_T[k] - \frac{TC}{V}$$

Proof: The proof is similar to that of Theorem 2 and is omitted for brevity. □

Combining the results of the above two theorems gives the following performance guarantee for all positive integers T and K:

$$\begin{aligned}\phi(\overline{u}(KT)) \quad \geq \quad & \frac{1}{K}\sum_{k=0}^{K-1}\psi_T[k] \\ & -\frac{TC}{V} - M\sqrt{\frac{2C + 2V\phi^{max}}{KT}}\end{aligned}$$

Rearranging and taking a limit as $K \to \infty$ gives:

$$\liminf_{K\to\infty}\left[\phi(\overline{u}(KT)) - \frac{1}{K}\sum_{k=0}^{K-1}\psi_T[k]\right] \geq -\frac{TC}{V}$$

For any $T > 0$, the right-hand-side above can be made arbitrarily small by an appropriately large value of V, with a corresponding convergence time tradeoff as specified in part (b) of Theorem 3.

VI. More conservative constraints

This section considers a variation that requires the achieved utility of each player *on each round* to be at least as large as the corresponding baseline utility for that round. That is, the time average constraint (14) is replaced by the more restrictive constraint:

$$u_i(t) \geq x_i(t) \ \forall i \in \mathcal{N}\ , \ \forall t \in \{0, 1, 2, \ldots\} \tag{32}$$

where we recall that:

$$\begin{aligned}u_i(t) &= \hat{u}_i(\alpha(t), \omega(t)) \\ x_i(t) &= \hat{u}_i(b(t), \omega(t))\end{aligned}$$

The resulting problem of interest is:

Maximize:

$$\liminf_{t\to\infty}\phi(\overline{u}_1(t), \ldots, \overline{u}_N(t)) \tag{33}$$

Subject to:

$$\hat{u}_i(\alpha(t), \omega(t)) \geq \hat{u}_i(b(t), \omega(t)) \forall i \in \mathcal{N}, \forall t \in \{0, 1, 2, \ldots\} \tag{34}$$

$$\alpha_i(t) \in \mathcal{A}_i \ \forall i \in \mathcal{N}, \forall t \in \{0, 1, 2, \ldots\} \tag{35}$$

where $\phi(u_1, \ldots, u_N)$ is again assumed to be a concave and continuous function over $\prod_{i=1}^{N}[0, u_i^{max}]$. The above problem is always feasible because the decisions $\alpha(t) = b(t)$ for all $t \in \{0, 1, 2, \ldots\}$ trivially satisfy the constraints (34)-(35). If the constraints of the above problem are satisfied, then the constraints of problem (13)-(15) are also satisfied (but not vice versa). Enforcing more restrictive constraints can reduce the optimal objective function value. However, it provides players an immediate guarantee that the suggested decisions are at

least as good as the baseline decisions, whereas problem (13)-(15) provides a similar guarantee only in the limit of a time average over multiple rounds.

The above problem can be written more simply as follows. Define:

$$\begin{aligned} u(t) &= (u_1(t), \ldots, u_N(t)) \\ \overline{u}(t) &= (\overline{u}_1(t), \ldots, \overline{u}_N(t)) \end{aligned}$$

For each $b \in \mathcal{A}_1 \times \cdots \times \mathcal{A}_N$ and each $\omega \in \Omega$, define $\mathcal{A}(b, \omega)$ as the set of all action vectors $\alpha = (\alpha_1, \ldots, \alpha_N)$ that satisfy:

$$\begin{aligned} \hat{u}_i(\alpha, \omega) \geq \hat{u}_i(b, \omega) \quad \forall i \in \mathcal{N} \\ \alpha_i \in \mathcal{A}_i \quad \forall i \in \mathcal{N} \end{aligned}$$

The set $\mathcal{A}(b, \omega)$ is non-empty for all (b, ω) because it contains the element b. The problem (33)-(35) is equivalent to the following:

$$\text{Maximize:} \quad \liminf_{t\to\infty} \phi(\overline{u}(t)) \tag{36}$$

$$\text{Subject to:} \quad \alpha(t) \in \mathcal{A}(b(t), \omega(t)) \ \forall t \in \{0, 1, 2, \ldots\} \tag{37}$$

If the set $\mathcal{A}(b(t), \omega(t))$ contains only the single element $b(t)$ for all $t \in \{0, 1, 2, \ldots\}$, then there are no decisions and the game manager is forced to choose $\alpha(t) = b(t)$ for all rounds t. In this case, the problem is so restricted that the game manager cannot provide any utility gain. However, for many problems the sets $\mathcal{A}(b(t), \omega(t))$ can have more than one element.

A. Weighted sum of utilities

First consider the special case:

$$\phi(u_1, \ldots, u_N) = \sum_{i=1}^{N} \theta_i u_i$$

where θ_i are real numbers. In this special case, the following simple strategy is optimal: Every round $t \in \{0, 1, 2, \ldots\}$, the game manager observes $b(t)$ and $\omega(t)$ and then chooses a vector $\alpha(t) \in \mathcal{A}(b(t), \omega(t))$ to maximize the following expression:

$$\sum_{i=1}^{N} \theta_i \hat{u}_i(\alpha(t), \omega(t)) \tag{38}$$

To see why this is optimal, consider any sequences $\{b(t)\}_{t=0}^{\infty}$ and $\{\omega(t)\}_{t=0}^{\infty}$. A sequence of actions $\{\alpha(t)\}_{t=0}^{\infty}$ is said to be a *feasible sequence of actions* if $\alpha(t) \in \mathcal{A}(b(t), \omega(t))$ for all t. Define $\{\alpha^*(t)\}_{t=0}^{\infty}$ as a feasible sequence of actions that maximize (38) on every round t. Let $\{\alpha'(t)\}_{t=0}^{\infty}$ be any alternative feasible sequence of actions. Then for all rounds $\tau \in \{0, 1, 2, \ldots\}$ one has:

$$\sum_{i=1}^{N} \theta_i \hat{u}_i(\alpha^*(\tau), \omega(\tau)) \geq \sum_{i=1}^{N} \theta_i \hat{u}_i(\alpha'(\tau), \omega(\tau))$$

Fix t as a positive integer. Summing the above over $\tau \in \{0, 1, 2, \ldots, t-1\}$ and dividing by t proves that:

$$\sum_{i=1}^{N} \theta_i \overline{u}_i^*(t) \geq \sum_{i=1}^{N} \theta_i \overline{u}_i'(t) \tag{39}$$

where $\overline{u}_i^*(t)$ is the time average utility of player i over the first t rounds under the actions $\{\alpha^*(\tau)\}_{\tau=0}^{\infty}$, while $\overline{u}_i'(t)$ is the corresponding time average utility under the actions $\{\alpha'(\tau)\}_{\tau=0}^{\infty}$. The inequality (39) is true for all rounds $t \in \{1, 2, 3, \ldots\}$ and so it is also true when taking limits as $t \to \infty$.

B. General concave function of utilities

A naive attempt to solve (36)-(37) might consider the policy of observing $b(t)$ and $\omega(t)$ every round t and then choosing $\alpha(t) \in \mathcal{A}(b(t), \omega(t))$ to maximize $\phi(u_1(t), \ldots, u_N(t))$. The previous subsection shows this naive policy is optimal in the special case when $\phi(\cdot)$ is linear. However, it is not necessarily optimal when $\phi(\cdot)$ is concave but nonlinear.

The problem (36)-(37) is similar to an *opportunistic scheduling problem* for wireless networks, as considered in [13][14][12] via the drift-plus-penalty approach and in [15][16][17][18] via different approaches. The works [13][14][12] transform a problem involving the maximization of a concave function of time averages into a problem of maximizing the time average of a function. The same approach is fruitful in this game theory context. Define proxy variables $\gamma_i(t)$ for all $i \in \mathcal{N}$. Define $\gamma(t) = (\gamma_1(t), \ldots, \gamma_N(t))$. Consider the following problem:

$$\text{Max:} \quad \liminf_{t\to\infty} \tfrac{1}{t} \textstyle\sum_{\tau=0}^{t-1} \phi(\gamma(t)) \tag{40}$$

$$\text{Subj to:} \quad \lim_{t\to\infty} [\overline{\gamma}_i(t) - \overline{u}_i(t)] = 0 \ \forall i \in \mathcal{N} \tag{41}$$

$$0 \leq \gamma_i(t) \leq u_i^{max} \ \forall i \in \mathcal{N}, \forall t \in \{0, 1, 2, \ldots\} \tag{42}$$

$$\alpha(t) \in \mathcal{A}(b(t), \omega(t)) \ \forall t \in \{0, 1, 2, \ldots\} \tag{43}$$

Lemma 4: If $\{\alpha(t)\}_{t=0}^{\infty}$ and $\{\gamma(t)\}_{t=0}^{\infty}$ are sequences of decisions that satisfy the constraints (41)-(43), then $\{\alpha(t)\}_{t=0}^{\infty}$ is a feasible sequence of control actions for the original problem (36)-(37) with the following utility guarantee:

$$\liminf_{t\to\infty} \phi(\overline{u}(t)) \geq \liminf_{t\to\infty} \frac{1}{t} \sum_{\tau=0}^{t-1} \phi(\gamma(\tau)) \tag{44}$$

Proof: Let $\{\alpha(t)\}_{t=0}^{\infty}$ and $\{\gamma(t)\}_{t=0}^{\infty}$ be sequences that satisfy (41)-(43). Fix $t > 0$. By Jensen's inequality and concavity of the $\phi(\gamma)$ function one has:

$$\begin{aligned} \frac{1}{t} \sum_{\tau=0}^{t-1} \phi(\gamma(\tau)) &\leq \phi(\overline{\gamma}(t)) \\ &\leq \phi(\overline{u}(t)) + M||\overline{\gamma}(t) - \overline{u}(t)|| \end{aligned} \tag{45}$$

where M is the Lipschitz constant for the function $\phi(\cdot)$. Taking the $\liminf$ of both sides and using (41) proves (44). Furthermore, since $\alpha(t)$ satisfies (43), it is a feasible sequence of control actions for the original problem. □

Lemma 4 suggests that one should make actions in an effort to solve the problem (40)-(43). Any decisions that are feasible for the problem (40)-(43) and that produce a "large" value of the objective function will also be feasible for the original problem with a corresponding objective function value that is at least as large. Let v_1^{opt} and v_2^{opt} be the supremum objective function values for the problems (36)-(37) and (40)-(43), respectively. The above lemma implies that $v_1^{opt} \geq v_2^{opt}$. If $(b(t), \omega(t))$ is ergodic, it turns out that $v_1^{opt} = v_2^{opt}$ with

probability 1 (see [12]), although this is not necessarily true for general non-ergodic problems.

The Lyapunov optimization method can be used to treat the problem (40)-(43). To enforce the constraints (41), for each $i \in \mathcal{N}$ define a virtual queue:

$$Z_i(t+1) = Z_i(t) + \gamma_i(t) - u_i(t) \tag{46}$$

Define $Z(t) = (Z_1(t), \ldots, Z_N(t))$. Define $L(t) = \frac{1}{2}||Z(t)||^2$ and $\Delta(t) = L(t+1) - L(t)$. As before, it can be shown that:

$$\begin{aligned}&\Delta(t) - V\phi(\gamma(t))\\ &\leq D - V\phi(\gamma(t))\\ &+ \sum_{i=1}^{N} Z_i(t)[\gamma_i(t) - \hat{u}_i(\alpha(t), \omega(t))]\end{aligned} \tag{47}$$

where D is a constant. Minimizing the right-hand-side of (47) every round t results in the following algorithm: Every round t, the game manager observes $Z(t)$, $b(t)$, $\omega(t)$. Then:

- (Proxy variables) Choose $\gamma(t) = (\gamma_1(t), \ldots, \gamma_N(t))$ as the solution to:

$$\begin{aligned}\text{Maximize:} \quad & V\phi(\gamma(t)) - \textstyle\sum_{i=1}^{N} Z_i(t)\gamma_i(t)\\ \text{Subject to:} \quad & 0 \leq \gamma_i(t) \leq u_i^{max} \ \ \forall i \in \mathcal{N}\end{aligned}$$

- (Suggestions) Choose $\alpha(t) = (\alpha_1(t), \ldots, \alpha_N(t))$ as the solution to:

$$\text{Maximize:} \quad \textstyle\sum_{i=1}^{N} Z_i(t)\hat{u}_i(\alpha(t), \omega(t)) \tag{48}$$

$$\text{Subject to:} \quad \alpha(t) \in \mathcal{A}(b(t), \omega(t)) \tag{49}$$

  Then send these suggestions to the corresponding players.
- (Queue update) Update $Z_i(t)$ for $i \in \mathcal{N}$ via (46).

For simplicity, it is assumed throughout that for all $t \in \{0, 1, 2, \ldots\}$ there exists an $\alpha(t)$ that solves (48)-(49) (else, the C-additive approximation theory of [12] can be used). It follows that the resulting $\alpha(t)$ sequence is a feasible sequence of control actions for the original problem (36)-(37). The next subsection analyzes its performance.

C. Analysis for conservative constraints and concave $\phi(\cdot)$

Fix sequences $\{\omega(t)\}_{t=0}^{\infty}$ and $\{b(t)\}_{t=0}^{\infty}$. Fix a positive integer T. For $k \in \{0, 1, 2, \ldots\}$ consider the following T-slot lookahead problem, which uses decision variables $\gamma = (\gamma_1, \ldots, \gamma_N)$ and $\alpha(t) = (\alpha_1(t), \ldots, \alpha_N(t))$:

Maximize:

$$\phi(\gamma) \tag{50}$$

Subject to:

$$\gamma_i = \frac{1}{T}\sum_{\tau=kT}^{kT+T-1} \hat{u}_i(\alpha(\tau), \omega(\tau)) \ \ \forall i \in \mathcal{N} \tag{51}$$

$$\alpha(\tau) \in \mathcal{A}(b(\tau), \omega(\tau)) \ \ \forall \tau \in \{kT, \ldots, (k+1)T - 1\} \tag{52}$$

Define $\psi_T[k]$ as the supremum objective function value in the above problem. Thus, for any $\epsilon > 0$, there is a sequence of decisions $\alpha(\tau)$ for $\tau \in \{kT, \ldots, (k+1)T - 1\}$ and a vector $\gamma = (\gamma_1, \ldots, \gamma_N)$ that together satisfy (51)-(52) and also satisfy:

$$\psi_T[k] - \epsilon \leq \phi(\gamma) \leq \psi_T[k] \tag{53}$$

Theorem 5: Fix $V \geq 0$ and assume the algorithm in the previous subsection is used with this V and with $Z_i(0) = 0$ for all $i \in \mathcal{N}$. For all positive integers T and K the following holds:

$$\phi(\overline{u}(KT)) \geq \frac{1}{K}\sum_{k=0}^{K-1} \psi_T[k] - \frac{DT}{V} - M\sqrt{\frac{2D + 2V\phi^{max}}{KT}}$$

where M is the Lipschitz constant for the function $\phi(\cdot)$ and D is the constant in (47). In particular, for all positive integers T one has:

$$\liminf_{t\to\infty} \phi(\overline{u}(t)) \geq \liminf_{K\to\infty} \frac{1}{K}\sum_{k=0}^{K-1} \psi_T[k] - \frac{DT}{V} \tag{54}$$

The theorem is proven in three parts.

Proof: (Part 1) This part proves that:

$$\phi(\overline{u}(KT)) \geq \frac{1}{KT}\sum_{\tau=0}^{KT-1} \phi(\gamma(\tau)) - M\sqrt{\frac{2(D + V\phi^{max})}{KT}} \tag{55}$$

To this end, note that (47) implies that for all rounds τ:

$$\begin{aligned}&\Delta(\tau) - V\phi(\gamma(\tau))\\ &\leq D - V\phi(\gamma^*(\tau))\\ &+ \sum_{i=1}^{N} Z_i(\tau)[\gamma_i^*(\tau) - \hat{u}_i(\alpha^*(\tau), \omega(\tau))]\end{aligned} \tag{56}$$

where $\gamma^*(\tau)$ and $\alpha^*(\tau)$ are any vectors that satisfy $\alpha(\tau) \in \mathcal{A}(b(\tau), \omega(\tau))$ and $\gamma_i^*(\tau) \in [0, u_i^{max}]$ for all $i \in \mathcal{N}$. Choose $\alpha^*(\tau) = b(\tau)$ and $\gamma^*(\tau) = (\gamma_1^*(\tau), \ldots, \gamma_N^*(\tau))$ where $\gamma_i^*(\tau) = \hat{u}_i(b(\tau), \omega(\tau))$ for all $i \in \mathcal{N}$. Substituting these choices into (56) gives:

$$\Delta(\tau) - V\phi(\gamma(\tau)) \leq D - V\phi(\gamma^*(\tau))$$

and hence:

$$\Delta(\tau) \leq D + V\phi^{max}$$

Summing over $\tau \in \{0, \ldots, KT - 1\}$ gives:

$$\frac{1}{2}||Z(KT)||^2 - \frac{1}{2}||Z(0)||^2 \leq KT(D + V\phi^{max})$$

Rearranging terms and using $||Z(0)|| = 0$ gives:

$$\frac{||Z(KT)||}{(KT)} \leq \sqrt{\frac{2(D + V\phi^{max})}{KT}}$$

From (46) it holds that and $||Z(KT)||/t = ||\overline{\gamma}(KT) - \overline{u}(KT)||$ and so:

$$||\overline{\gamma}(KT) - \overline{u}(KT)|| \leq \sqrt{\frac{2(D + V\phi^{max})}{KT}}$$

By (45) it follows that:

$$\frac{1}{KT}\sum_{\tau=0}^{KT-1} \phi(\gamma(\tau)) \leq \phi(\overline{u}(KT)) + M\sqrt{\frac{2(D + V\phi^{max})}{KT}}$$

which proves (55). □

Proof: (Part 2) This part shows that:

$$\frac{1}{KT}\sum_{\tau=0}^{KT-1} \phi(\gamma(\tau)) \geq \frac{1}{K}\sum_{k=0}^{K-1} \psi_T[k] - \frac{DT}{V} \tag{57}$$

To this end, note that summing (56) over $\tau \in \{kT, \ldots, (k+1)T-1\}$ gives:

$$\begin{aligned}
&\sum_{\tau=kT}^{(k+1)T-1} \Delta(\tau) - V \sum_{\tau=kT}^{(k+1)T-1} \phi(\gamma(\tau)) \\
&\leq DT - V \sum_{\tau=kT}^{(k+1)T-1} \phi(\gamma^*(\tau)) \\
&+ \sum_{\tau=kT}^{(k+1)T-1} \sum_{i=1}^{N} Z_i(\tau)[\gamma_i^*(\tau) - \hat{u}_i(\alpha^*(\tau), \omega(\tau))] \\
&\leq DT^2 - V \sum_{\tau=kT}^{(k+1)T-1} \phi(\gamma^*(\tau)) \\
&+ \sum_{i=1}^{N} Z_i(kT) \sum_{\tau=kT}^{(k+1)T-1} [\gamma_i^*(\tau) - \hat{u}_i(\alpha^*(\tau), \omega(\tau))]
\end{aligned}$$

where the final step is similar to a step in the proof of Theorem 2. Fix $\epsilon > 0$ and define $\gamma^* = (\gamma_1^*, \ldots, \gamma_N^*)$ and $\alpha^*(\tau)$ for $\tau \in \{kT, \ldots, (k+1)T-1\}$ as the vectors that satisfy (51), (52), (53), and define $\gamma_i^*(\tau) = \gamma_i^*$ for all $\tau \in \{kT, \ldots, (k+1)T-1\}$. Substituting these into the above inequality gives:

$$\begin{aligned}
&\sum_{\tau=kT}^{(k+1)T-1} \Delta(\tau) - V \sum_{\tau=kT}^{(k+1)T-1} \phi(\gamma(\tau)) \\
&\leq DT^2 - VT\psi_T[k] + VT\epsilon
\end{aligned}$$

Taking $\epsilon \to 0$ and then summing over $k \in \{0, \ldots, K-1\}$ gives:

$$L(KT) - L(0) - V \sum_{\tau=0}^{KT-1} \phi(\gamma(\tau)) \leq DT^2K - VT \sum_{k=0}^{K-1} \psi_T[k]$$

Dividing by VKT and using the fact that $L(KT) - L(0) \geq 0$ gives:

$$-\frac{1}{KT} \sum_{\tau=0}^{KT-1} \phi(\gamma(\tau)) \leq \frac{DT}{V} - \frac{1}{K} \sum_{k=0}^{K-1} \psi_T[k]$$

which proves (57). □

Proof: (Part 3) This part proves (54). To this end, note that parts 1 and 2 together imply:

$$\phi(\overline{u}(KT)) \geq \frac{1}{K} \sum_{k=0}^{K-1} \psi_T[k] - \frac{DT}{V} - M\sqrt{\frac{2D + 2V\phi^{max}}{KT}}$$

Taking a $\liminf$ of both sides as $K \to \infty$ gives:

$$\liminf_{K\to\infty} \phi(\overline{u}(KT)) \geq \liminf_{K\to\infty} \frac{1}{K} \sum_{k=0}^{K-1} \psi_T[k] - \frac{DT}{V} \qquad (58)$$

It remains to show that the left-hand-side of (58) can be replaced by $\liminf_{t\to\infty} \phi(\overline{u}(t))$. To do this, fix t as a positive integer. Let K_t be the non-negative integer such that $K_tT \leq t < (K_t+1)T$. Then:

$$\overline{u}(t) = \overline{u}(K_tT)\frac{K_tT}{t} + \frac{\sum_{\tau=K_tT}^{t-1} u(\tau)}{t}$$

In particular:

$$\overline{u}(t) = \overline{u}(K_tT) - \overline{u}(K_tT)\frac{(t - K_tT)}{t} + \frac{\sum_{\tau=K_tT}^{t-1} u(\tau)}{t}$$

Thus:

$$||\overline{u}(t) - \overline{u}(K_tT)|| \leq \frac{\sqrt{\sum_{i=1}^{N}(T-1)u_i^{max}}}{t}$$

Thus:

$$\phi(\overline{u}(t)) \geq \phi(\overline{u}(K_tT)) - M\frac{\sqrt{\sum_{i=1}^{N}(T-1)u_i^{max}}}{t}$$

Taking a $\liminf$ of both sides as $t \to \infty$ gives:

$$\begin{aligned}
\liminf_{t\to\infty} \phi(\overline{u}(t)) &\geq \liminf_{t\to\infty} \phi(\overline{u}(K_tT)) \\
&\geq \liminf_{K\to\infty} \phi(\overline{u}(KT))
\end{aligned}$$

This together with (58) proves the result. □

VII. Conclusion

This paper considers a stochastic repeated game where players share information and a baseline decision with a game manager at the beginning of each round. The manager provides suggestions that, if taken, maximize a concave function of average utilities across players subject to the constraint that each player receives a time average utility at least as good as it would get if all players used their baseline strategies. A more conservative scenario was also considered where the utility guarantee is enforced every round, rather than in a time average. A Lyapunov optimization algorithm was developed that satisfies the constraints and that ensures the concave function of utilities is close to (or better than) an average of T-slot lookahead values that are computed with knowledge of T rounds into the future, regardless of the sample path. This shows that a simple causal algorithm can achieve a target defined in terms of future knowledge.

References

[1] D. P. Foster and R. V. Vohra. Calibrated learning and correlated equilibrium. *Games and Economic Behavior*, vol. 21, pp. 40-55, 1997.

[2] S. Hart and A. Mas-Colell. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, vol. 68, no. 5, pp. 1127-1150, Sept. 2000.

[3] D. Fudenberg and D. K. Levine. Conditional universal consistency. *Games and Economic Behavior*, vol. 29, no. 1-2, pp. 104-130, Oct. 1999.

[4] E. Solan and N. Vieille. Correlated equilibrium in stochastic games. *Games and Economic Behavior*, vol. 38, pp. 362-399, 2002.

[5] M. J. Neely. A Lyapunov optimization approach to repeated stochastic games. *ArXiv technical report,arXiv:1310.2648v3*, Feb. 2014.

[6] J. F. Nash. Non-cooperative games. *Annals of Mathematics*, vol. 54, pp. 286-295, 1951.

[7] J. F. Nash. Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences of the United States of America*, vol. 36, pp. 48-49, 1950.

[8] R. Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, vol. 1, pp. 67-96, 1974.

[9] R. Aumann. Correlated equilibrium as an expression of Bayesian rationality. *Econometrica*, vol. 55, pp. 1-18, 1987.

[10] H. Moulin and J. P. Vial. Strategically zero-sum games: The class of games whose completely mixed equilibria cannot be improved upon. *International Journal of Game Theory*, vol. 7, no. 3/4, pp. 201-221, 1978.

[11] M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, Cambridge, MA, 1994.
[12] M. J. Neely. *Stochastic Network Optimization with Application to Communication and Queueing Systems*. Morgan & Claypool, 2010.
[13] M. J. Neely, E. Modiano, and C. Li. Fairness and optimal stochastic control for heterogeneous networks. *IEEE/ACM Transactions on Networking*, vol. 16, no. 2, pp. 396-409, April 2008.
[14] L. Georgiadis, M. J. Neely, and L. Tassiulas. Resource allocation and cross-layer control in wireless networks. *Foundations and Trends in Networking*, vol. 1, no. 1, pp. 1-149, 2006.
[15] A. Stolyar. Maximizing queueing network utility subject to stability: Greedy primal-dual algorithm. *Queueing Systems*, vol. 50, no. 4, pp. 401-457, 2005.
[16] A. Eryilmaz and R. Srikant. Fair resource allocation in wireless networks using queue-length-based scheduling and congestion control. *IEEE/ACM Transactions on Networking*, vol. 15, no. 6, pp. 1333-1344, Dec. 2007.
[17] R. Agrawal and V. Subramanian. Optimality of certain channel aware scheduling policies. *Proc. 40th Annual Allerton Conf. on Communication, Control, and Computing, Monticello, IL*, Oct. 2002.
[18] H. Kushner and P. Whiting. Asymptotic properties of proportional-fair sharing algorithms. *Proc. 40th Annual Allerton Conf. on Communication, Control, and Computing, Monticello, IL*, Oct. 2002.